Designing Co-Creative AI for Public Spaces

Duri Long
Georgia Institute of Technology
Atlanta, GA, USA
duri@gatech.edu

Mikhail Jacob
Georgia Institute of Technology
Atlanta, GA, USA
mikhail.jacob@gatech.edu

Brian Magerko
Georgia Institute of Technology
Atlanta, GA, USA
magerko@gatech.edu

ABSTRACT

Artificial intelligence (AI) is becoming increasingly pervasive in our everyday lives. There are consequently many common misconceptions about what AI is, what it is capable of, and how it works. Compounding the issue, opportunities to learn about AI are often limited to audiences who already have access to and knowledge about technology. Increasing access to AI in public spaces has the potential to broaden public AI literacy, and experiences involving co-creative (i.e. collaboratively creative) AI are particularly well-suited for engaging a broad range of participants. This paper explores how to design co-creative AI for public interaction spaces, drawing both on existing literature and our own experiences designing co-creative AI for public venues. It presents a set of design principles that can aid others in the development of co-creative AI for public spaces as well as guide future research agendas.

Author Keywords

collaboration; public displays; reflection on design processes; co-creative AI; human-centered AI

CCS Concepts

•**Human-centered computing** → *Empirical studies in interaction design;* •**Computing methodologies** → *Artificial intelligence;*

INTRODUCTION

Artificial intelligence (AI) is becoming increasingly prevalent in our everyday lives–in places as personal as our social media news feeds, cars, and homes. However, there are still many misconceptions regarding what exactly AI is, what it is capable of, and how it works. There are few opportunities to learn about AI outside of a university setting, and the online resources that do exist primarily cater to an already tech-savvy audience (e.g. online tutorials, cloud services for machine learning (ML), ML APIs).

Recent exploratory research has investigated how to communicate "big ideas" about AI to a non-expert audience–for example, the idea that computers perceive the world through sensors or that computers learn from data [80]. Many tool-kits that cater to non-expert audiences focus on providing learners with the tools to "tinker" with AI through the programming of robots or AI-powered cloud services such as Cozmo, Sphero, Alexa, and Google Assistant (e.g. [81, 17, 41]), see [80] for an exhaustive list). These tool-kits have (to date) been designed primarily for K-12 classroom activities and are accompanied by curricula and worksheets (e.g. [81]). Museums and other public spaces can serve as alternative venues for AI literacy initiatives, complementing interventions in the formal education sphere and broadening access to opportunities to both interact with and learn about AI to both adults and children who may not have AI devices in their homes or schools.

Co-creative (i.e. collaboratively creative) AI (c.f. [16, 56, 32, 38, 93]) may be particularly well-suited for AI literacy initiatives in public spaces. *Co-creation* (i.e. "a social creativity process 'leading to the emergence and sharing of creative activities and meaning in a socio-technical environment"' [42, 22]) has been shown to be a powerful in-road for learning about computing-related disciplines, even for communities that may not otherwise be interested in–or feel included in–computing [57, 7, 28]. The social, open-ended nature of co-creative experiences also makes them well-suited for engaging participants in free-choice informal learning environments [20]. AI systems that can engage as active participants in co-creative activities therefore have potential to serve as avenues for AI literacy in public spaces.

Recent work has called on AI researchers to share their work with the public [80], but there is not much guidance provided on how to effectively transform a research project into a public installation. In this paper, we aim to contribute to a better understanding of our central research question: *what design principles should be considered when designing co-creative AI for public interaction spaces in order to better educate and engage the public?* We will address this question by examining the broader literature on designing technology for public spaces as well as engaging in reflective practice [75, 95] to draw on knowledge gained from our own experiences developing co-creative AI for public spaces.

RELATED WORK

Research that touches on how to design co-creative AI for use in public spaces is still in its early stages. There is a body of research focused on how to design and develop expressive [58] and co-creative [42] AI. Some of these projects have been used in public spaces (e.g. [60]), although the literature focuses

C&C '19, June 23–26, 2019, San Diego, CA, USA

ACM 978-1-4503-5917-7/19/06. . . $15.00
DOI: https://doi.org/10.1145/3325480.3325504

mostly on designing for co-creativity and artistic expression, not public use. Recent literature at the intersection of HCI and AI/ML research is investigating new design considerations that may need to be taken into account when developing AI/ML algorithms for (mostly online) public use [92, 62]. There is also a significant body of research on interactive technology and its use in public art installations and museum exhibits that can inform our work. For example, Snibbe and Raffle present a number of principles for designing social immersive media in museums, and while they do not focus specifically on AI, a number of their projects involve AI technologies [78]. In this section, we explore six different themes/common threads that we have identified in the literature.

Design for Apprehendability

Public-facing installations are installed in places such as museums where visitors frequently suffer from cognitive overload [19, 1], finding it hard to focus due to an excess of external stimuli. Allen suggests that exhibits should be designed such that people interacting with them for the first time can immediately understand their "purpose, scope, and properties" [1], thus reducing cognitive load. She calls this *immediate apprehendability*. Snibbe and Raffle [78] reinforce this idea, emphasizing the importance of designing exhibits that are immediately responsive to participants' actions. Others have echoed this notion by calling for more naturalistic, apprehendable modes of interaction such as Kinect-based embodied interaction in art museums [44], tangible programming blocks in science museums [33], and voice interaction with a robot tour guide at a public expo [40]. When designed with object affordances [25, 64] and intuitive mappings [2] in mind, such interfaces are able to engage a wide range of participants, including those with low computational literacy [34].

Design for Collaboration

Public spaces are also social spaces. Most people visiting museums come in groups [30], and research in urban design has revealed that facilitating social interaction is a critical component to consider when designing public spaces in cities [89]. Social interaction and collaboration also play an important role in the learning process [86, 49]. However, depending on their design, introducing interactive technologies in public spaces can actually inhibit social interaction [30]. A variety of existing work suggests that technology in public spaces should be designed explicitly to be supportive of group interaction (e.g. [33, 78]). Snibbe and Raffle even argue that "user engagement...should become richer as more people interact" [78]. Strategies like providing multiple access points to an exhibit [34], allowing for multiple levels of engagement [11], and making interactions with the exhibit/artifact visible to passersby [91] can all help to encourage collaboration. Bengler et al. emphasize that even if "active collaboration" is not occurring due to contextual reasons (e.g. crowding or a lack of familiarity between participants), exhibits should at least allow for participants to feel connected to each other [3].

Design for Learning

Engagement with public exhibits is often very brief (typical "dwell time" at many science museums is around 30 seconds [35]), but time-on-task is an important factor in learning environments [23]. Designers of public exhibits focused on facilitating learning often advocate designs that encourage "extended and repeated engagement" [34] and "active prolonged engagement" [36]. Quality of engagement also plays an important role in learning–Humphrey et al. suggest that learning is best facilitated by exhibits that encourage active sense-making [36] and Hornecker and Sifter suggest that exhibits for learning should facilitate intellectual engagement that extends beyond just free play [34]. In addition to designing the exhibit itself to facilitate active learning, others have augmented exhibits with interactive "contemplation rooms" that strive to provide additional information to promote learning about the exhibit content [44].

One way to encourage people to learn about AI is to create "explainable" agents that can help participants to understand their decision-making process [47]. Recent work has shown that increasing transparency about both the existence and capabilities of AI algorithms can change users' experiences and perceptions of AI [18]. Explainability of agent reasoning is easier to achieve when the agent's task is well-specified and goal-oriented [67, 48]. Since co-creative tasks are often poorly-specified and exploratory (or autotelic), co-creative agent explainability is challenging. AI that uses more "black-box" techniques like neural networks and reinforcement learning in place of symbolic reasoning makes explainability even more difficult to accomplish. Notable advances have been made in interpreting the results of neural networks that operate on input data like images [65, 94]. Others have taken inspiration from how humans justify their decision-making post-hoc and use separate AI subsystems to generate explanations after an agent has already acted (e.g. [18]).

Design for Creative Engagement

Designing to facilitate creative engagement has also been a focus of a number of public art installations, and can contribute to an understanding of how to design co-creative agents for similar spaces. Bilda et al. suggest that participants engaging with public interactive artworks progress sequentially from being initially attracted to an installation to engaging in sustained interaction with the installation and finally to relating their interaction with the installation to experiences outside of the installation space [4]. Designers should consider this progression and design features of the installation that encourage participants to move from one stage to the next [4].

Wouters et al. also suggest a number of design principles for encouraging participants to transition between different interaction roles (e.g. passer-by, participant, dropout) when interacting with public art installations [91]. These principles include advertising the installation and making it visible to passers-by, creating spaces that facilitate observation and social interaction, incorporating collaborative features and surprising moments into the design, allowing for multiple degrees of interaction with the system, and creating an activation loop in which dropouts can encourage passers-by or audience members to become participants.

Finally, a body of research has contributed to a better understanding of what factors increase perceptions of creativity in interactions with creativity support tools and co-creative agents. A computational system's perceived level of skill, imagination, and appreciation of its own outputs directly affect collaborator perceptions of the system's creativity in turn [12]. Additionally, the ability of a computational system to pleasantly surprise its collaborator over the course of an interaction is an important factor in facilitating co-creative interactions [26]. A participant's perceived sense of control, challenge, and satisfaction [3] during a computer-mediated creative interaction also encourage participants to enter a flow state [13]. Designing for increased immersion, satisfaction, enjoyment, collaboration, exploration, and expressiveness [8] also allows computational tools to better support participant creativity.

Design for Robustness and Safety

Introducing technological artifacts in public spaces requires a level of robustness and safety that is not typical of many research projects. Jensen et al. discuss the safety features built into a robotic tour guide, placing an emphasis on the need for redundant features and emergency safety controls [40]. Horn et al. also highlight inexpensiveness and reliability as key design principles to consider when creating tangible user interfaces for museums [33]. It is also important to consider the interplay between the environment and the technology. For example, both Jensen et al. and Bengler et al. discuss dealing with technological complications due to ambient noise/light in public spaces [40, 3].

Balancing Design Concerns with Research Goals

Balancing the aforementioned design principles for public spaces with research and evaluation needs is not always straightforward. Designing for public spaces presents unique challenges for researchers, such as difficulty comparing iterations of a project across different venues, conflict between "the artistic impulse to improve an exhibit...with the need for experimental control", aligning research questions with the theme of the creative work, and obtaining informed consent in public spaces [68].

In addition to these concerns, understanding how to assess and rigorously evaluate interactive technology in public spaces is still an open research question. Heath and von Lehm call for new methods of evaluating interactions at the "exhibit face" [30], and recent work has explored a variety of strategies for evaluating/understanding factors such as learning, collaboration, and creativity including mixed-method approaches [3], video analysis [14, 36], and conversation analysis [70].

Takeaways

The body of literature on designing technology for museums and art spaces provides a foundation for exploring the design of co-creative agents for public spaces. However, none of this work focuses specifically on introducing co-creative agents into public spaces, and only a few of the papers presented draw on experiences from more than one installation. We will build on the existing work in the next sections by reflecting on our experience developing a variety of co-creative agents for a diverse range of public spaces over the course of several years to suggest design principles that both support and add to the principles that have been highlighted in the literature so far.

METHODOLOGY - REFLECTIVE PRACTICE

Schön originally characterized design as a reflective practice, in which designers reflect on their actions in order to contribute to a growing body of design knowledge and methodologies [75]. Reflective practice has since been recognized within the interaction design and HCI communities as a useful research methodology in which designed artifacts can become "exemplars" or "conduit[s] for research findings to easily transfer to the HCI research and practice communities" [95], and reflection on the design process and the themes that the artifacts embody can improve both design practice [95] and contribute to design theory [24]. Snibbe and Raffle's paper on designing social immersive media for museums is a good example of reflective design practice utilized in a relevant domain [78].

In this paper, we reflect on our practice as designers of co-creative AI in public spaces, drawing on our experience iteratively designing a variety of different co-creative installations involving AI technologies over the past five years. We define *public space* as any physical space to which the general public has access–including urban spaces, museums, art galleries, and public events or "happenings" [43] (although our experiences draw primarily on the latter three). For the sake of this paper, online/virtual "public" spaces are not considered. The exhibits we will be discussing in this paper are briefly described below.

LuminAI (Fig 1) is an installation in which a human participant and an AI agent can dance together [38]. A Kinect motion capture device is used to detect the participant's motion, which is visualized as a virtual "shadow" on a projection screen. Next to the shadow is a humanoid "agent", which dances by analyzing the participant's movement and responding with a movement that it deems to be similar in terms of parameters such as energy, tempo, or size (adapted from Viewpoints movement theory [6]). The agent interactively learns gestures from the participant as they dance together.

The *Robot Improv Circus* (Fig 1) is an interactive installation in which participants can take turns improvising open-ended comedic interactions using abstract props together with a robotic scene partner in virtual reality (VR), while an audience can view and reflect on the improvised performance from outside the circus tent through portals into the virtual world. The virtual improviser generates action candidates based on the physical attributes of the prop and the ongoing interaction. It chooses its response in real-time in order to satisfy an intrinsic motivation or drive to follow a given "creative arc" over the duration of the improvised performance. The creative arc is a designer-specified trajectory through the space of novelty, unexpectedness, and quality, where these properties are computed for each action that the agent considers. Creative arcs allow the designer some artistic control over the agent's improvisation while still providing participants with a qualitatively evolving experience over the course of the improvisation.

Sound Happening (Fig 1) is a playspace in which participants can collaboratively create music together by moving colorful balls around a defined interaction space [54]. An overhead camera and computer vision software tracks the location and the color of the balls and a computer generates music accord-

Figure 1: Installations, from left to right. 1) LuminAI; 2) The audience view of the Robot Improv Circus VR experience. Text in the robot's speech bubble reads, "I am looking with my kaleidoscope"; 3) Sound Happening; 4) The Shape of Story

ing to those parameters. Certain events (e.g. all balls meeting in the center of the interaction space) trigger special sound effects.

Shape of Story (Fig 1) is a story circle experience in which participants collectively create a story line-by-line [53]. Participants speak into a listening device and AI in narrative understanding is used to interpret their words. The semantic meaning of participants' words is translated into a symbolic visual language, which is drawn in real-time on a projected "painting" in the center of the story circle. The result is a narrative art piece that is co-created by the participants telling the story and the agent creating the visualization. The physical installation surrounding the story circle is set up to create a comfortable, intimate [52] environment, not unlike gathering around a campfire to share stories. Participants enter and exit the installation through a hallway where they can view "paintings" created by previous groups. The interaction experience is led by a facilitator.

We have publicly exhibited the aforementioned projects at a variety of different venues–including academic conferences, art festivals, museums, and other local venues. We have experience working with a variety of installation timeframes (month, week, and day-long), participants of all ages from all over the world, and different sets of spatial and environmental constraints. We have also worked with a variety of partners during the design process, including museum educators and practitioners, artists, evaluators, and other researchers. A complete list of the public installations of our projects used to develop the design principles described in this paper is included in Table 1. We also provide a list of which design principles were derived from which installation in Table 1.

In keeping with reflective practice in the context of design research as our methodology, we reflected on a) the physical, technical, and experience design of the installations and how they changed over time; b) our design process and research/evaluation methodologies; and c) the nature of participant interaction with the installations. Data that we consulted in our reflection process included: group discussions and archived meeting notes detailing our collective experiences during the design, evaluation, and research processes; observations and findings from video analysis, interviews, and field notes from previous studies of participant interaction with the installations (see the following papers for detailed descriptions of methodologies for each of these studies: [55, 54, 53]; installations in which formal research studies were conducted are marked with an asterisk in Table 1); and video and images of different iterations of the designed artifacts.

Some questions we considered when reflecting on our design practice and developing the design principles outlined in the remainder of this paper are listed here: How did the technical/physical/experience design change over time? What aspects of the design worked well in public settings? What aspects failed or led to unexpected challenges? Was there anything surprising about the ways in which participants interacted in installations? Did participants collaborate with each other or with the AI agent(s) in our designed systems? Did our designed artifacts facilitate creative interaction? If so, which aspects of the installations contributed most effectively to the creative interaction? How did participants make sense of the AI system(s) they were interacting with? What feedback/advice did we receive from our research partners/collaborators and participants? Which aspects of our development, research design, and evaluation processes were successful and which aspects needed improvement? Were any of our observations supported by the existing literature?

DESIGN PRINCIPLES

This section presents a set of design principles we have derived from reflecting on our experiences installing co-creative AI installations in a variety of public spaces. These design principles are broken down into three overarching categories: 1) technology design; 2) interaction design; and 3) research design. Within each of these categories, we further organize the design principles in sub-categories, which are visualized in Figure 4 and described in detail in the remainder of this section.

Technology Design

We first explore aspects of technical design that can contribute to the successful installation of co-creative AI in a public space, including issues relating to the maintenance and adaptability of the system and cognitive capabilities of the AI agent that contribute to the participant experience.

Maintenance

DP1 (Modularity): The software architecture of the system should be adaptable and modular, enabling easy replacement of components that can be quickly and reliably modified based

Project	Date	Event/Location	Users	Data Collected	Design Principles
LuminAI	2013	The Window Project, Atlanta*	50		1, 2, 3, 9, 10, 11, 12
LuminAI	2014	The Art + Science Museum, Singapore	150		1, 2, 3, 9, 10, 11, 12
LuminAI	2015	Creativity and Cognition, Glasgow	35		1, 2, 3, 9, 10, 11, 12
LuminAI	2016	Field Experiment, Atlanta*	100	Video, interviews	1, 2, 3, 5, 7, 8, 9, 10, 11, 12, 13, 16
LuminAI	2017	ACC Creativity & Innovation Festival, D.C.	300		2, 3, 4, 6, 7, 8, 11, 16
LuminAI	2017	Children's Museum of Pittsburgh*	150	Video	1, 2,, 4, 6, 7, 8, 11, 12, 13, 14, 15, 16
LuminAI	2018	CODAME Art+Tech Festival, SF	200		5, 9, 10, 11, 15, 16
RIC	2018	TechRec, Atlanta*	120	Surveys, interviews	3, 4, 6, 8, 9, 10, 11, 13, 16, 17
RIC	2018	The Biltmore, Atlanta*	75	Interviews	3, 4, 5, 6, 8, 9, 10, 11, 13, 16, 17
SH	2017	Clough Art Crawl, Atlanta	75		11
SH	2017	Children's Museum of Pittsburgh*	165	Video	1, 2, 7, 8, 11, 12, 13, 14, 15, 16, 17
SH	2019	ACC Creativity & Innovation Festival, D.C.	30		1, 9, 11
SoS	2017	Eyedrum Art and Music Gallery, Atlanta	20		8, 9, 11, 14

Table 1: Installation summary, chronologically by project. Abbreviations: the Robot Improv Circus (RIC), Sound Happening (SH), the Shape of Story (SoS). Informal observations were noted at all installations, even when other data was not collected. * indicates that a formal research study was conducted. All user counts are estimated.

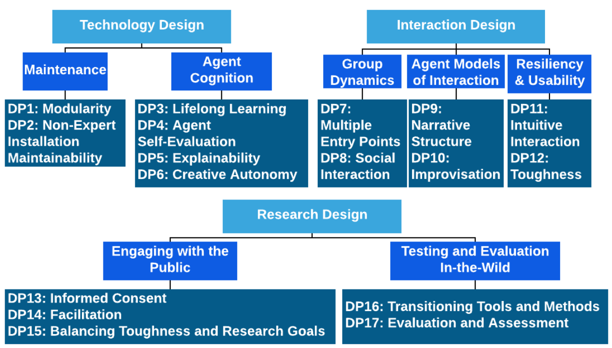

Figure 2: Design principles summary chart

on the installation setting/needs. Designing for modularity makes the need to customize exhibits for different physical, creative or technical venue constraints (a need identified in [68] and our own work) much easier. For example, *LuminAI* has been used for dance performances as well as installations, engaging both adults and children depending on the venue. Each of these settings has necessitated different technical capabilities, and a modular code base enables rapid adaptation, such as the ability to quickly swap out the Kinect for a more precise MoCap suit, rapidly change the look-and-feel of the avatar and its surroundings, or change the movement theory the agent uses to reason.

DP2 (Non-Expert Installation Maintainability): AI systems for public spaces should be developed such that non-experts like museum staff members can easily run/restart the software and modify certain aspects of agent decision making. The installation should also be accompanied by good documentation. When exhibiting interactive installations there are several options for reacting to failure modes or unexpected/dynamic changes in the installation context. One option is to attempt to automatically monitor the performance of an installation for undesirable changes in behavior–such as reduced agent reactivity to participant inputs–in order to reboot the agent or run troubleshooting scripts. However, as we discovered in late 2013 running *LuminAI* for over two weeks projected on a storefront, where the public could interact with the virtual agent as they walked past, attempting to predict all the different failure modes and continuously monitoring the system for all of them proved to be resource-prohibitive both in system resources and developer time.

Remote human technical support can be provided for interactive installations as an alternative to automatic performance monitoring (as we did for *LuminAI* in Glasgow in 2015). This is a solution that can potentially work well for large research teams or creative agencies that can afford to offer dedicated technical support services since a technical expert can address installation issues in a timely manner without disclosing intellectual property. However, remote technical support can be particularly challenging for research groups that operate under constrained resources and have to prioritize other aspects of the research process over post-installation technical support. The need for coordination between technical experts and venue staff across time zones to ensure that the remote installation is accessible, connected to the internet, and any fixes are tested can also be challenging.

While exhibiting both *LuminAI* and *Sound Happening* at the Children's Museum of Pittsburgh in 2017, we realized that if we wanted to install our projects as permanent exhibits, a practical solution to handling installation failure modes was to ensure that on-the-floor staff members could easily run/restart the project and make minor adjustments to it as needed throughout the day. Documentation and built-in tools that allow non-experts to reboot the project, troubleshoot errors, and make certain customizations (e.g. adjust agent decision-making parameters or swap out sound samples) can serve as a practical solution to this problem.

Agent Cognition

*DP3 (Lifelong Learning): Co-creative AI agents that learn from collaborators over the course of their run-time should be able to process their growing knowledge/experience base *without slowing down over time.* A co-creative AI agent requires the ability to access and learn from a large dataset or knowledge-base about the world and experiences within it in order to be able to match the breadth and depth of creativity that people offer the agent and expect to receive in return. Many embodied creative domains like dance, pretend play, and theater are notable for their lack of large-scale, diverse, annotated datasets, since embodied knowledge or data is time-consuming and expensive to collect. Agents capable of lifelong learning (c.f. [77]) can remedy this problem since they can learn interactively from the human participants themselves. Increasing database size can be a concern with installations in public settings, particularly if only the local computer's finite memory capacity is used. Cloud storage provides an alternative to finite local storage but is dependent on reliable high-speed network bandwidth which is not present in many public spaces.

A larger problem than finite storage space for a growing knowledge-base is that of needing to efficiently search and retrieve knowledge or data from the knowledge-base [87]. The literature on case-based reasoning, lifelong machine learning, and interactive learning present some strategies that can facilitate searching a growing database without sacrificing speed. One such strategy is to forget older, less useful, less-frequently-used, or bad data [37]. Unfortunately, this can lead to missing data in learned patterns without exacting a pattern management or correction overhead. Abstraction can be used to generalize data points into representative prototypes that can be compared against in lieu of individual data points. However, this can also lead to challenges maintaining learned patterns in the data if the prototypes change over time due to lifelong learning. Certain search strategies are particularly well-suited for searching large knowledge-bases. Branch and bound search [50] can be used to aggressively prune the agent's search space as it looks for a suitable solution. Alternatively, a stochastic satisficing search can be conducted until a data point that is "just good enough" is found by the agent or time expires [39].

DP 4 (Agent Self-Evaluation): Co-creative AI agents that learn from collaborators and generate creative offers should be able to evaluate and distinguish between "good" and "bad" percepts and actions. Distinguishing between percepts based on quality enables the agent to both improve the quality of its responses and save storage space for higher-quality data to learn from, making lifelong learning more feasible. This is especially important in public spaces where non-experts may provide the agent with a lot of low-quality or repetitive data. Some obviously "bad" data can be avoided by minimizing opportunities for undesirable interactions. For instance, in *LuminAI*, we disable detection of finger movement to eliminate the possibility of recording lewd gestures, and we limit the size and setup of the interaction space so participants cannot sit in front of the installation (thereby teaching the agent how to sit instead of dance). Agents can also use heuristics and function approximation to evaluate to what degree new data points are novel, unexpected, or high quality (for a given definition of quality) (c.f. [46]). These evaluative models could potentially be trained on or guided by feedback from participants for personalization or corrected by experts (e.g. [16]).

Computational models for evaluating the novelty, unexpectedness, quality, or aesthetics of an agent's potential responses help the agent to explicitly select responses with a specific degree of those qualities during a collaboration, instead of relying purely on tacit procedural knowledge to implicitly attempt this task [85, 84]. This is particularly helpful within open-ended co-creative AI installations where there is not a single clear goal, objective function, or other reward function for which to optimize the agent's behavior. For example, in the *Robot Improv Circus* [39], agent self-evaluation is used to explicitly consider whether a potential response fits a creative arc that the agent is trying to follow over the course of the improvisation. This also allows the agent to display feedback to human participants about what degree of novelty, unexpectedness, and quality it perceived for their input actions. Possessing the capability to computationally evaluate the novelty, unexpectedness, and quality of responses thus provides the agent with a vocabulary for dialogue with the human participant about the ongoing co-creative partnership regardless of the accuracy of those evaluative models.

While there are numerous computational models for evaluating novelty, unexpectedness, and quality (c.f. [46]), designers of co-creative experiences should consider the following constraints when choosing a specific set of models: the perspective from which the model performs the evaluation (agent, human, or audience), the dynamism (i.e. change over time) of the model as the agent's experience increases, the role of feedback in the model, and expertise of the agent relative to its collaborator (c.f. [39]).

DP 5 (Explainability): Consider making AI systems "explainable" by making their decision-making process transparent, incorporating supplementary interactions to help participants understand the inner-workings of the agent, and considering how interface design plays a role in understandability. Providing access to AI systems in public spaces is one step towards promoting widespread AI literacy, but access does not necessarily lead to increased understanding. For example, in previous studies of *LuminAI* in public spaces, many participants vastly over or under-estimated the capabilities of the AI agent they were interacting with, and some did not even understand that they were engaging with AI [38, 55]–a finding that is consistent with prior research on human-AI interaction [88]. How can we claim that the human and AI are co-creating if the human participant does not have an accurate theory of mind representing how the agent works?

We have begun to investigate how to make the agents in our projects more "explainable", but we are still in the early stages of this aspect of our research and much remains for future work. Some strategies that we have tried draw on research in explainable AI (see Related Work). We have explored the use of short text snippets in both *LuminAI* and the *Robot Improv Circus* as a way of communicating aspects of agent reasoning that are not otherwise apparent. For instance, in earlier versions of *LuminAI*, the agent simply waved at the participant to initiate an interaction. Now, text displayed above the agent says "Will you teach me how to dance, human?". Informally, we have observed that many more participants verbally recognize that the agent is able to learn from them when the text prompt is visible.

Research has shown that the specific framing of an explanation also impacts the perceived creativity of an act [9]. For example, a technical explanation stating that a response was selected because it was "15% more novel than the other responses in the decision space" would be perceived as being more creative if it were to be re-formulated into the sentence, "I was really curious about this action that I hadn't tried before and wanted to try it out to see if it worked in this situation." A notable component of the explanation is some level of fictitious intentionality and agency. We have observed this effect in our own work–placing speech bubbles around agent statements in the *Robot Improv Circus* caused participants to perceive intentionality, stating that they thought the agent was trying to do "something specific".

Finally, we have found that other factors (beyond text explanations) can also play a role in participant interpretations of the system. For instance, in one version of *LuminAI*, we found that the agent visualization (a cloud of moving fireflies) was aesthetically pleasing but obscured some of the agent's motions and made it difficult for participants to understand what was going on [55]. We later transitioned to using a more humanoid avatar for the dancer.

DP 6 (Creative Autonomy): Allow for modulating the AI agent's creative autonomy. Generally, we have found that the level of creative autonomy that is appropriate for co-creative AI agents is heavily dependent on context. Human collaborators tend to want more control over the behavior and actions of co-creative AI when they are using the AI as a performance technology or creativity support tool [15, 72]. For installations or exhibits designed for non-expert audiences, it can be beneficial to give agents more creative autonomy so that they can provide non-experts with scaffolding for the co-creative task and better support the creative collaboration [15]. However, even with novices, we have found that it is important for people to see how their actions influence the co-creative process, as perceived control plays an important role in helping people to achieve creative "flow" [3, 59]. Our investigation of leader-follower roles [90] for the agent within the *LuminAI* installation showed that there was a qualitative difference in participant experiences between an agent that preferred to follow most of the time and one that would actively switch between leading and following. The novice participants seemed to prefer the agent that followed their actions more. Others have also found that humans dislike when agents take the lead in interactions [45], a phenomenon that could (at least in the case of *LuminAI*) be explained by the chameleon effect [10].

One solution to the need for variable creative autonomy is to design the agent architecture so that the degree of creative autonomy can be modulated according to the context of the collaboration. In the *Robot Improv Circus*, the agent architecture is explicitly driven to select actions according to a designer-specified creative arc for the agent's performance over the course of the improvised performance. This provides the experience designer with a level of creative control over the degree of novelty, unexpectedness, and quality of the agent's actions during the co-creation. The respective creative arcs for the agent could be varied to control the agent's creative autonomy based on its use as a performance technology or as an installation in an informal learning context.

Interaction Design

This section explores design considerations related to how people interact with co-creative agents in public spaces. We discuss factors including group dynamics, agent models of interaction, and resilience and usability.

Group Dynamics

DP 7 (Multiple Entry Points): Design for multiple entry points and levels of engagement (even when you are targeting a particular audience). Unlike in a controlled classroom or laboratory environment, we cannot determine who is going to interact with AI in public spaces. How do we make installations engaging and informative for people with different age levels and levels of expertise? This is especially important to consider when designing for family group interactions in museums.

We design our installations with multiple entry points for participants of different ages (as suggested in [78]). For instance, in *Sound Happening*, even very young children can (and love to) push around balls and hear the noise they make, but the interaction potential is deep enough that some adults experiment with making coordinated rhythms with the balls for long periods of time [54]. In *LuminAI*, young children often just enjoy dancing with their Kinect "shadows", whereas older children and adults recognize and take into account the agent's motions. In both projects, participants are able to creatively express themselves in a variety of ways.

DP 8 (Social Interaction): Design agents that can interact socially with humans and/or help to facilitate human-human social interaction. As mentioned in Related Work, public spaces are social spaces (75% of visitors to museums come in groups [30]), and there is a relationship between learning and social interaction [86, 49]. However, most co-creative AI are designed for one-on-one interaction (e.g. [16, 56, 61, 51]), and interactive museum exhibits involving technology have more broadly been critiqued for discouraging social interaction [30]. Equipping agents with social intelligence–such as the ability to take turns dynamically [90]–can help to make human-agent social interaction more naturalistic.

Designers should also consider how installations can support human collaboration (a quality that has repeatedly been identified as being an important component of public installations [91, 33, 55]). Even something as simple as providing a way for multiple group members to peripherally interact with an exhibit at once (rather than just passively observing as audience members) can transform the nature of social interaction. For example, in *LuminAI*, we display up to three people's Kinect "shadows" on the screen (even though the agent only responded to the movements of the participant standing closest to the screen). We saw a wide variety of social interactions as a result of enabling the additional shadows, including a couple salsa dancing, two friends coordinating a synchronized dance performance, and a group of teenagers performing in a dance circle. Similarly, in *Sound Happening*, having multiple balls available to interact with led to the emergence of parent-child joint play [54]. In the *Robot Improv Circus*, the audience can watch someone interacting in VR through video portals outside the circus tent and can provide feedback using their voice and gestures to trigger in-game reward systems. This led to several groups of friends encouraging each other to try different actions with the props.

The physical design of the space can also contribute to promoting social interaction–in *LuminAI*, we found that a partially enclosed dome provided a semi-private dance floor where people felt more comfortable interacting than in front of a two-dimensional screen [55]. In *Shape of Story*, we explicitly designed an enclosed storytelling space to encourage participants to share stories with each other (and the system) that they may not feel comfortable divulging to a large audience.

Agent Models of Interaction

DP 9 (Narrative Structure): Consider the narrative trajectory of the participant's full experience before, during, and after interaction with the installation. While there are many ways for participants to interact with public installations and there is no one "right way", providing some narrative structure to the experience can encourage participants to engage more deeply [11, 4, 78]. Designers should consider the following questions:

What does the initial interaction look like? A variety of literature has pointed to the importance of "attractors" or "advertisements" that invite participants to interact with public installations [91, 4, 52, 33]. In our work, we have experimented with agent actions that welcome participants into the installation (e.g. a waving or dancing agent in *LuminAI*), initial text prompts that provide contextual or welcoming information (*LuminAI*, the *Robot Improv Circus*), and pre-interaction demonstrations or performances to entice passerby to join in (*LuminAI*, *Sound Happening*). We have also explored how designing physical entrance areas and facilitated "welcoming" stages [52] can help prepare participants for interaction (*Shape of Story* [53]).

Is there a narrative structure to the collaborative interaction? The structure of the body of the interaction is also important to consider. The *LuminAI* agent is capable of human-agent turn-taking–this social leading/following behavior provides some structure to the body of the interaction [90]. We have also started to incorporate the beginnings of narrative sequencing into co-creative agent "turns"–for instance, having an agent walk to a prop and pick it up prior to playing a generated action in the *Robot Improv Circus*, or playing a sequence of related gestures that build on each other in *LuminAI*. We have also experimented with text prompts and signs that encourage participants to try different actions (a strategy also employed by [36, 76]). Facilitated guidance can also serve as a way of creating a narrative structure. For example, in *Shape of Story*, a facilitator guided participants through the experience of telling a collaborative story together [53].

How does the interaction end? Everyone leaves an interaction eventually, and the way that an interaction ends can play a role in what the participant takes away from the experience. "Designing for dropout" [91] can encourage personal reflection and social sharing of experiences at the end of an interaction [52, 4, 44]. While some of our projects involve very simple conclusions (such as an agent waving goodbye), in others we have explored how to facilitate reflection. For example, in *Shape of Story*, we encouraged participants leaving the exhibit to reflect on the meaning of the story that was visualized by looking at example stories posted in the hallway outside [53].

DP 10 (Improvisation): Equipping agents with knowledge of how to improvise, or creatively respond to novel situations, can help to make the agent more socially interactive (DP 8), naturalistic (DP 11), and "tough" (DP 12). In-the-wild interactions with co-creative agents can often yield unexpected situations and novel inputs. One way to deal with this, especially in the context of co-creative interaction, is utilizing improvisational techniques. A variety of work has studied how humans improvise together in different domains like jazz [73, 66, 31, 5], theater [74, 56], and dance [83, 69, 79]. Many of our co-creative AI projects have used improvisational techniques to improve the quality of interaction and creative response [38, 39, 56]. For instance, trading fours/trading eights is a technique used in jazz improvisation in which musicians alternate playing four or eight bars of music, "continu[ing] in the spirit or mood established by the prior players, responding to, and building on, the prior musician's eight bars" [74]. We use a similar technique in *LuminAI*, in which the agent, upon recognizing a gesture, responds with a series of three gestures–one that mimics the human gesture, one that transforms it in some way, and one novel gesture that is similar to the observed gesture in some way (e.g. rhythm, tempo, size). In the *Robot Improv Circus* as well, the agent uses a similar set of strategies adapted from human improvisational practice in order to prioritize its search of potential responses to follow a particular creative arc over the course of the performance.

Resilience and Usability

DP 11 ("Intuitive" Interaction): Design AI agents that can interact in naturalistic, embodied, and/or culturally recognizable ways with people. People quickly abandon exhibits that are difficult to interact with in public spaces [1, 78]. Interaction modalities that require high computational literacy can also discourage certain groups (e.g. the elderly) from participating [34]. Designing to mitigate these issues is referred to in the literature as designing for apprehendability [1, 33] or responsiveness [78] (see Related Work).

All of our installations involving co-creative AI employ embodied or naturalistic interaction methods. We use a variety of technology to achieve this, including the Microsoft Kinect depth sensor (*LuminAI*), the HTC Vive room-scale VR system (the *Robot Improv Circus*), microphones and voice recognition (*Shape of Story*), and webcam color tracking with Max MSP (*Sound Happening*). Cultural understanding of socially familiar activities such as ball play can also make interaction easier for participants [33, 78], a technique that we leveraged in *Sound Happening* to encourage adult-child play [54].

DP 12 ("Toughness"): Design AI systems and their associated installations for "toughness"; in other words, design AI installations such that they can withstand long-term heavy *use by many people.* In public spaces, AI installations may be interacted with over long periods of time, in unexpected ways, and possibly in a rough manner. While many technology developers strive to create robust systems, the level of "toughness" needed in public spaces is easy to underestimate and can involve factors that are not often considered in adult-only research environments (e.g. avoiding choking hazards, preventing theft of installation materials, preparing the installation for rough interactions such as throwing, kicking, hitting). In addition to designing for robustness and safety (see Related Work), toughness means:

The installation does not break when interacted with over long periods of time (see DP 3). In addition to technical implementation and physical robustness, this includes considering interference issues related to the environment. As noted in Related Work, Jensen et al. and Bengler et al. both discuss noise and light interference in public spaces, issues that we have also dealt with in our projects that involve projection and/or sound [40, 3]. Other contextual issues to consider include WiFi connectivity/dependability, sensor interference, and the amount and type of foot traffic that your installation location receives. Designers (in collaboration with museum partners or artists) should also take into consideration what level of maintenance is required to keep the installation running (see DP 2).

Designers should try to consider all of the possible (ideal and less than ideal) ways someone might interact with the exhibit and design such that none of those "break" the exhibit. This necessitates an iterative design research process [24], as it is impossible to foresee all possible use cases. Through iterative testing that we conducted as part of the Children's Museum of Pittsburgh's "Tough Art" program, we identified a need for and built a box to protect the Kinect from prying fingers in *LuminAI* and devised a solution to contain the balls within the interaction space for *Sound Happening*. Transportability, space considerations, and ease of set-up/take-down should also be considered during the design process. We found that installing *LuminAI* in a large geodesic dome encouraged social interaction [55], but the dome was difficult to transport/set up and too large to fit in most exhibit spaces. We are currently exploring a solution involving more mobile and compact scaffolding. Finally, designers should consider how social norms and object affordances play into the durability of the installation. For instance, in *Sound Happening*, we found that bouncy balls were quickly weaponized by kids who threw them at each other and the walls. When replaced with balloons, which are not able to bounce or be thrown far, theft became an issue as balloons were perceived as being "free". Beach balls ended up being a workable final solution [54].

Research Design

This section explores how designing AI for public spaces relates to research agendas and goals. We discuss issues including challenges associated with engaging with the public and testing/evaluating exhibits in-the-wild.

Engaging with the Public

DP 13 (Informed Consent): It can be difficult to obtain informed consent to study (and in particular, video record) participants' interactions with AI installations in public spaces. Beyond just getting IRB approval, designers should more deeply examine whether participants (especially children) know that they are participating in research when interacting with public exhibits and what that participation entails. Visitors should be reminded of consent as they enter the specific exhibit (not just upon entering an event or venue). This is often achieved via signage, but it is important to recognize the role that sign size and placement plays in determining whether participants realize they are participating in research. We have found that large signs placed at all entry points to an installation (and inside of the installation, if applicable) are most effective, a finding that is supported by other museum research [27].

The age of participants and the context of the venue can also affect one's ability to obtain informed consent–young children who are not being supervised might run into the exhibit without their parents. We found obtaining informed consent to be more straightforward at museums that were less free-play oriented for this reason. In addition, it is important to consider that school groups and friends often visit museums together, meaning that it cannot be assumed that all children's parents are present. Access-controlled spaces (with a facilitator present to enforce age limits and parental consent) may be necessary in some scenarios. In museum spaces, designated "research" areas or "living laboratories" [63] where researchers often study new installations can help make the line between "exhibit" and "research exhibit" clearer for the visitor.

DP 14 (Facilitation): Designers should consider different facilitation needs for different contexts. Some facilitation can help push interactions to "the next level", but heavy-handed facilitation can disrupt the exploratory nature of co-creative interaction. In informal learning spaces, facilitators that ask questions that promote further exploration can promote participant learning and engagement, but heavy-handed explanations might actually discourage further engagement. Recent research on facilitation suggests that it may be more productive for facilitators to make provocative comments or questions intended to push participants to the next level of learning/engagement (e.g. "Do you think you can teach the agent a dance move?"; "I wonder which ball controls which noise.") as opposed to approaching participants and offering extensive explanations or critiques [29].

Facilitation also plays an important role in art spaces. Loke and Khut [52] discuss how facilitators can play an important role in guiding participants through more intimate artistic experiences in public spaces. We explored this approach in our work with *Shape of Story*, where facilitators led participants through the entire interaction. Researchers also need to consider how the presence of facilitators relates to their research questions and evaluation instruments. If researchers are trying to measure the ability of the exhibit to foster in-the-wild learning or engagement, facilitation might skew results. Other research questions may be more compatible with facilitation.

DP 15 (Balancing "Toughness" and Research Goals): Researchers should recognize the "tough" nature of public exhibits. It can be difficult to balance the need for "tough" exhibits (DP 12) with research goals. There is a conflict between quickly iterating to test out new research ideas vs. taking the time to ensure robustness, modularity, good documentation, and easy non-expert maintainability (see DP 2). This can slow down the research process but is necessary for long-term public installation and work with partners who are concerned with durability. As a research community, it would be helpful to place value on the production of durable public installations involving co-creative AI (rather than only valuing novel prototypes of AI systems). Embedded in these artifacts is knowledge about designing for public interaction [24]. In addition, such artifacts allow us to explore novel research questions with a wider community of study participants.

Testing and Evaluation In-the-Wild

DP 16 (Transitioning Tools and Methods): Rapid prototyping (and the tools/research methods associated with it) is useful in the beginning/testing stages of a project, but for researchers concerned with developing "tough" installations, it is important to transition from prototyping platforms and research methods to industry-standard platforms and in-the-wild studies. We create initial builds of many of our projects using Processing (a great tool for rapid prototyping). However, Processing is difficult to use for creating large-scale software projects and industry standard game engines/editors like Unity3D provide better installation performance in the long-term. Yang points out that new tools developed specifically for designers and artists looking to incorporate AI in their work may aid in the design process, especially since rapid prototyping can be challenging when working with ML [92]. Examples of such tools include Fiebrink et al.'s Wekinator tool, designed to enable artists to use ML to control real-time performance [21], and van Allen's prototyping platform that allows designers to simulate AI devices with physical components prior to transitioning to an actual hardware prototype [82].

Transitioning from more structured workshop studies to in-the-wild observations is also critical. Findings from controlled studies often do not transfer to "the wild" [71], and museum workshops held during prototyping stages may not fully illuminate the social interactions that occur with an exhibit on-the-floor [30]. It is therefore important to test project iterations at various stages in-the-wild.

DP 17 (Evaluation and Assessment): Designers of co-creative AI should consider how to evaluate their installations in order to assess the degree to which they facilitate learning/creative expression, recognizing that new exploring new modes of evaluation and assessment may be necessary. Evaluating co-creative experiences in public spaces is difficult–what makes an interaction a "good" one? It is challenging to assess the degree to which an installation facilitates learning/meaning-making, creativity, and/or collaboration. Recent work has explored how to quantify co-creative interactions in order to better understand the quality of collaboration and sense-making [14]. There is also a body of research that explores how to evaluate the creative products generated by AI systems (see Related Work). We are currently drawing on literature from museum studies and studies of co-creative systems we have developed in order to better understand learning that occurs in co-creative interactions, from physical, social, emotional, and intellectual perspectives.

FUTURE WORK

A number of the design principles we listed identify areas for our current and future research. We plan on building on the *LuminAI* system to communicate some of the "big ideas" of AI [80], focusing on incorporating affordances that support participant learning (DP 5) and fostering social interaction (DP 8) in order to support family group learning. Within the *Robot Improv Circus* installation, we also plan on further exploring how to improve agent explainability (DP 5) to better support informal learning for the audience and creative autonomy (DP 6) to better support use as a performance technology. Finally, we are exploring new modes of evaluating co-creative experiences through qualitative analysis of our systems in use in-the-wild.

In addition, the design principles presented in this paper are intended for use when designing and developing AI for public spaces. However, for the most part they do not address the early idea-development stages of the design process. Recent research has called for designers to consider ML and AI more broadly as "design materials", or materials that can be used to develop new product forms [92]. We suggest that co-creative AI is particularly well-suited as a material for designing public installations, as the social and open-ended, improvisational nature of interaction present in most co-creative systems complements the free-choice, group style of interaction in public spaces. When generating ideas for novel installations involving co-creative AI, we draw inspiration from both the strengths of co-creative AI as a design material and the design principles presented in this paper, while focusing on how the co-creative AI can facilitate experiences that elicit creative exploration, surprise, awe, joy, play, and social and embodied interaction. We hope to further explore co-creative AI as a design material in future work.

CONCLUSION

In this paper, we have contributed a practice-based understanding of the challenges that come with designing co-creative AI for public spaces, and an initial look at some of the solutions that we have explored. This contribution can be valuable to other co-creativity researchers looking to share their work with the public as well as museum exhibitors and artists looking at how they might incorporate AI in museums and gallery spaces. We encourage other researchers to not only consider these design principles when developing co-creative AI for public spaces, but also to consider exploring some of the identified gaps in existing literature as areas for future research.

ACKNOWLEDGEMENTS

We would like to thank past and present members of the Expressive Machinery Lab, the Children's Museum of Pittsburgh, Eyedrum Art and Music Gallery, and the Digital Interactive Liberal Arts Center at Georgia Tech for their support of this research. The Robot Improv Circus project is funded by a Creative Curricular Initiatives grant from the Georgia Tech Office of the Arts and the Georgia Tech Arts Council.

REFERENCES

[1] Sue Allen. 2004. Designs for learning: Studying science museum exhibits that do more than entertain. *Science Education* 88, 1 (2004), S17.

[2] Alissa N Antle, Greg Corness, and Milena Droumeva. 2009. Human-computer-intuition? Exploring the cognitive basis for intuition in embodied interaction. *International Journal of Arts and Technology* 2, 3 (2009), 235–254.

[3] Ben Bengler and Nick Bryan-Kinns. 2014. In the wild: evaluating collaborative interactive musical experiences in public settings. In *Interactive experience in the digital age*. Springer, 169–186.

[4] Zafer Bilda, Ernest Edmonds, and Linda Candy. 2008. Designing for creative engagement. *Design Studies* 29, 6 (2008), 525–540.

[5] Margaret A. Boden. 2003. *The Creative Mind - Myths and Mechanisms (2. ed.)*. Routledge.

[6] Anne Bogart and Tina Landau. 2004. *The Viewpoints Book: A Practical Guide to Viewpoints and Composition*. Theatre Communications Group, New York : St. Paul, MN.

[7] Leah Buechley, Mike Eisenberg, Jaime Catchen, and Ali Crockett. 2008. The LilyPad Arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 423–432. DOI: http://dx.doi.org/10.1145/1357054.1357123

[8] Erin A. Carroll, Celine Latulipe, Richard Fung, and Michael Terry. 2009. Creativity factor evaluation: towards a standardized survey metric for creativity support. In *Proceedings of the seventh ACM conference on Creativity and cognition*. ACM, 127–136. http://dl.acm.org/citation.cfm?id=1640255

[9] John William Charnley, Alison Pease, and Simon Colton. 2012. On the Notion of Framing in Computational Creativity.. In *ICCC*. 77–81.

[10] Tanya L Chartrand and John A Bargh. 1999. The chameleon effect: the perception–behavior link and social interaction. *Journal of personality and social psychology* 76, 6 (1999), 893.

[11] Luigina Ciolfi and Liam Bannon. 2002. Designing Interactive Museum Exhibits: Enhancing visitor curiosity through augmented artefacts. In *Eleventh European Conference on Cognitive Ergonomics*. Citeseer.

[12] Simon Colton. 2008. Creativity Versus the Perception of Creativity in Computational Systems.. In *AAAI spring symposium: creative intelligent systems*, Vol. 8.

[13] Mihaly Csikszentmihalyi. 2008. *Flow : the psychology of optimal experience* (1st harper perennial modern classics ed. ed.). Harper Perennial, New York.

[14] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Brenda Lin, and Brian Magerko. 2017. Creative sense-making: Quantifying interaction dynamics in co-creation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 356–366.

[15] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, and Brian Magerko. 2016. Co-creative drawing agent with object recognition. In *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*.

[16] Nicholas Davis, Chih-PIn Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 196–207.

[17] Stefania Druga. 2018. *Cognimates: from coding to teaching machines*. Ph.D. Dissertation. MIT.

[18] Upol Ehsan, Pradyumna Tambwekar, Larry Chan, Brent Harrison, and Mark Riedl. 2019. Automated Rationale Generation: A Technique for Explainable AI and its Effects on Human Perceptions. *arXiv preprint arXiv:1901.03729* (2019).

[19] Gary W Evans. 1995. Learning and the physical environment. *Public institutions for personal learning: Establishing a research agenda* (1995), 119–126.

[20] John H Falk and Lynn D Dierking. 2000. *Learning from museums: Visitor experiences and the making of meaning*. Altamira Press.

[21] Rebecca Fiebrink and Perry R Cook. 2010. The Wekinator: a system for real-time, interactive machine learning in music. In *Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)(Utrecht)*.

[22] Gerhard Fischer, Elisa Giaccardi, Hal Eden, Masanori Sugimoto, and Yunwen Ye. 2005. Beyond binary choices: Integrating individual and social creativity. *International Journal of Human-Computer Studies* 63, 4-5 (2005), 482–512.

[23] Wayne C Fredrick and Herbert J Walberg. 1980. Learning as a function of time. *The journal of educational research* 73, 4 (1980), 183–194.

[24] William Gaver. 2012. What should we expect from research through design?. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 937–946.

[25] J.J. Gibson. The theory of affordances. In *Perceiving, acting, and knowing*, R.E. Shaw and J Bransford (Eds.). Erlbaum Associates, Hillsdale, N.J.

[26] Kazjon Grace and Mary Lou Maher. 2016. Surprise-Triggered Reformulation of Design Goals.. In *AAAI*. 3726–3732.

[27] Joshua P Gutwill. 2003. Gaining visitor consent for research II: Improving the posted-sign method. *Curator: The Museum Journal* 46, 2 (2003), 228–235.

[28] Mark Guzdial. 2013. Exploring Hypotheses about Media Computation. In *Proceedings of the Ninth Annual International ACM Conference on International Computing Education Research*. ACM, San Diego, CA, USA, 19–26.

[29] Danielle B Harlow, Ron Skinner, and Sean O'Brien. 2017. Roll It Wall: Developing a framework for evaluating practices of learning. In *Proceedings of the 7th Annual Conference on Creativity and Fabrication in Education*. ACM, 14.

[30] Christian Heath and Dirk Vom Lehn. 2008. Configuring'Interactivity' Enhancing Engagement in Science Centres and Museums. *Social Studies of Science* 38, 1 (2008), 63–91.

[31] Robert Hodson. 2007. *Interaction, improvisation, and interplay in jazz*. Routledge.

[32] Guy Hoffman and Gil Weinberg. 2010. Shimon: an interactive improvisational robotic marimba player. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*. ACM, 3097–3102.

[33] Michael S. Horn, Erin Treacy Solovey, and Robert JK Jacob. 2008. Tangible programming and informal science learning: making TUIs work for museums. In *Proceedings of the 7th international conference on Interaction design and children*. ACM, 194–201. http://dl.acm.org/citation.cfm?id=1463756

[34] Eva Hornecker and Matthias Stifter. 2006. Learning from interactive museum installations about interaction design for public settings. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments*. ACM, 135–142.

[35] Sherry Hsi and Holly Fait. 2005. RFID enhances visitors' museum experience at the Exploratorium. *Commun. ACM* 48, 9 (2005), 60–65.

[36] Thomas Humphrey, Joshua Gutwill, and The Exploratorium APE Team. 2005. *Fostering Active Prolonged Engagement: The Art of Creating APE Exhibits*. Routledge, Abingdon, UK.

[37] Mikhail Jacob. 2017. Towards Lifelong Interactive Learning For Open-ended Embodied Narrative Improvisation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 502–507.

[38] Mikhail Jacob, Gaëtan Coisne, Akshay Gupta, Ivan Sysoev, Gaurav Verma, and Brian Magerko. 2013. Viewpoints AI. In *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE '13)*. AAAI, Northeastern University, Boston, MA, USA.

[39] Mikhail Jacob and Brian Magerko. 2018. Creative arcs in improvised human-computer embodied performances. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*. ACM, 62.

[40] BjÃűrn Jensen, Nicola Tomatis, Laetitia Mayor, Andrzej Drygajlo, and Roland Siegwart. 2005. Robots meet humans: Interaction in public spaces. *IEEE Transactions on Industrial Electronics* 52, 6 (2005), 1530–1546.

[41] Ken Kahn and Niall Winters. 2017. Child-friendly programming interfaces to AI cloud services. In *European Conference on Technology Enhanced Learning*. Springer, 566–570.

[42] Anna Kantosalo, Jukka M Toivanen, Ping Xiao, and Hannu Toivonen. 2014. From Isolation to Involvement: Adapting Machine Creativity Software to Support Human-Computer Co-Creation. In *ICCC*. 1–7.

[43] Allan Kaprow and Mimsy Lee. 1966. On Happenings. *The Tulane Drama Review* 10, 4 (1966), 281–283.

[44] Karen Johanne Kortbek and Kaj Grønbæk. 2008. Communicating art through interactive technology: new approaches for interaction design in art museums. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*. ACM, 229–238.

[45] Hatice Kose-Bagci, Kerstin Dautenhahn, and Chrystopher L Nehaniv. 2008. Emergent dynamics of turn-taking interaction in drumming games with a humanoid robot. In *17th IEEE Int. Symp. on Robot and Human Interactive Communication 2008 (RO-MAN'08)*. IEEE.

[46] Carolyn Lamb, Daniel G Brown, and Charles LA Clarke. 2018. Evaluating Computational Creativity: An Interdisciplinary Tutorial. *ACM Computing Surveys (CSUR)* 51, 2 (2018), 28.

[47] H Chad Lane, Mark G Core, Michael Van Lent, Steve Solomon, and Dave Gomboc. 2005. *Explainable artificial intelligence for training and tutoring*. Technical Report. University of Southern California Institute for Creative Technologies.

[48] Pat Langley, Ben Meadows, Mohan Sridharan, and Dongkyu Choi. 2017. Explainable Agency for Intelligent Autonomous Systems.. In *AAAI*. 4762–4764.

[49] Jean Lave and Etienne Wenger. 1991. *Situated learning: Legitimate peripheral participation*. Cambridge university press.

[50] Eugene L Lawler and David E Wood. 1966. Branch-and-bound methods: A survey. *Operations research* 14, 4 (1966), 699–719.

[51] Antonios Liapis, Georgios N Yannakakis, and Julian Togelius. 2013. Sentient Sketchbook: Computer-aided game level authoring.. In *FDG*. 213–220.

[52] Lian Loke and George Poonkhin Khut. 2014. Intimate aesthetics and facilitated interaction. In *Interactive Experience in the Digital Age*. Springer, 91–108.

[53] Duri Long, Sanjana Gupta, Jessica Brooke Anderson, and Brian Magerko. 2017. The Shape of Story: A Semiotic Artistic Visualization of a Communal Storytelling Experience. (2017).

[54] Duri Long, Hannah Guthrie, and Brian Magerko. 2018. Don't steal my balloons: designing for musical adult-child ludic engagement. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*. ACM, 657–662.

[55] Duri Long, Mikhail Jacob, Nicholas Davis, and Brian Magerko. 2017. Designing for socially interactive systems. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 39–50.

[56] B. Magerko, C. DeLeon, and P. Dohogne. 2011. Digital Improvisational Theatre: Party Quirks. AAAI Press, Reykjavík, Iceland.

[57] B. Magerko, J. Freeman, T. McKlin, M. Reilly, E. Livingston, S. McCoid, and A. Crews-Brown. 2016. EarSketch: A STEAM-based Approach for Underrepresented Populations in High School Computer Science Education. *ACM Transactions on Computing Education* 16, 4 (2016).

[58] Michael Mateas. 2001. Expressive AI: A hybrid art and science practice. *Leonardo* 34, 2 (2001), 147–153.

[59] Michael Mateas and Kate Compton. 2015. Casual Creators. In *Proceedings of the Sixth International Conference on Computational Creativity*. 228–235.

[60] Michael Mateas, Steffi Domike, and Paul Vanouse. 1999. Terminal time: An ideologically-biased history machine. *AISB Quarterly, Special Issue on Creativity in the Arts and Sciences* 102 (1999), 36–43.

[61] Marek P Michalowski, Selma Sabanovic, and Hideki Kozima. 2007. A dancing robot for rhythmic social interaction. In *Human-Robot Interaction (HRI), 2007 2nd ACM/IEEE International Conference on*. IEEE, 89–96.

[62] MinKyung Lee Amit Elazari Bar On Eric Gilbert Karrie Karahalios Motahhare Eslami, Kristen Vaccaro. 2019. User Attitudes towards Algorithmic Opacity and Transparency in Online Reviewing Platforms. In *ACM CHI 2019*.

[63] Lizzie Muller and Ernest Edmonds. 2006. Living laboratories: making and curating interactive art. *SIGGRAPH 2006 Electronic Art and Animation Catalog* (2006), 160–163.

[64] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Constellation.

[65] Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. 2018. The building blocks of interpretability. *Distill* 3, 3 (2018), e10.

[66] Jeff Pressing. 1988. Improvisation: methods and models. In *Generative processes in music*, John Sloboda (Ed.). Oxford University Press, Canada, 129–178.

[67] GabriÃńlle Ras, Marcel van Gerven, and Pim Haselager. 2018. Explanation methods in deep learning: Users, values, concerns and challenges. In *Explainable and Interpretable Models in Computer Vision and Machine Learning*. Springer, 19–36.

[68] Derek Reilly, Fanny Chevalier, and Dustin Freeman. 2014. Blending art events and HCI research. In *Interactive Experience in the Digital Age*. Springer, 153–168.

[69] Mônica m Ribeiro and Agar Fonseca. 2011. The empathy and the structuring sharing modes of movement sequences in the improvisation of contemporary dance. *Research in Dance Education* 12, 2 (2011), 71–85.

[70] Jessica Roberts and Leilah Lyons. 2017. Scoring Qualitative Informal Learning Dialogue: The SQuILD Method for Measuring Museum Learning Talk. Philadelphia, PA: International Society of the Learning Sciences.

[71] Yvonne Rogers. 2012. Interaction design gone wild: striving for wild theory. *Interactions* 18, 4 (2012), 58–62. `http://dl.acm.org/citation.cfm?id=1978834`

[72] Sarah Rubidge. 2002. Digital Technology in Choreography: Issues and Implications. In *Published in the Proceedings of the 17th Annual Symposium of the Dance Society of Korea, Seoul, Korea*.

[73] R Keith Sawyer. 2006. Group creativity: Musical performance and collaboration. *Psychology of Music* 34, 2 (2006), 148–165.

[74] R Keith Sawyer and Stacy DeZutter. 2009. Distributed creativity: How collective creations emerge from collaboration. *Psychology of aesthetics, creativity, and the arts* 3, 2 (2009), 81.

[75] Donald A Schön. 2017. *The reflective practitioner: How professionals think in action*. Routledge.

[76] Stephanie Shine and Teresa Y Acosta. 2000. Parent-Child Social Play in a Children's Museum. *Family Relations* 49, 1 (2000), 45–52.

[77] Daniel L Silver, Qiang Yang, and Lianghao Li. 2013. Lifelong Machine Learning Systems: Beyond Learning Algorithms.. In *AAAI Spring Symposium: Lifelong Machine Learning*, Vol. 13. 05.

[78] Scott S Snibbe and Hayes S Raffle. 2009. Social immersive media: pursuing best practices for multi-user interactive camera/projector exhibits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1447–1456.

[79] Carlota Torrents, Marta Castañer, Mária Dinušová, and M Teresa Anguera. 2010. Discovering new ways of moving: Observational analysis of motor creativity while dancing contact improvisation and the influence of the partner. *The Journal of Creative Behavior* 44, 1 (2010), 53–69.

[80] David Touretzky, Christina Gardner-McCune, Fred Martin, and Deborah Seehorn. 2019. Envisioning AI for K-12: What should every child know about AI?

[81] David S Touretzky. 2017. Computational thinking and mental models: From Kodu to Calypso. In *Blocks and Beyond Workshop (B&B), 2017 IEEE*. IEEE, 71–78.

[82] Philip van Allen. 2018. Prototyping Ways of Prototyping AI. *Interactions* 25, 6 (Oct. 2018), 46–51. DOI: http://dx.doi.org/10.1145/3274566

[83] Floor van Alphen. 2014. Tango and enactivism: first steps in exploring the dynamics and experience of interaction. *Integrative Psychological and Behavioral Science* 48, 3 (2014), 322–331.

[84] Dan Ventura. 2016. Mere generation: Essential barometer or dated concept. In *Proceedings of the Seventh International Conference on Computational Creativity*. Sony CSL, Paris, 17–24.

[85] Dan Ventura. 2017. How to build a CC system. In *Proceedings of the 8th International Conference on Computational Creativity, Atlanta*.

[86] Lev Semenovich Vygotsky. 1980. *Mind in society: The development of higher psychological processes*. Harvard university press.

[87] Christian Wagner. 2006. Breaking the knowledge acquisition bottleneck through conversational knowledge management. *Information Resources Management Journal (IRMJ)* 19, 1 (2006), 70–83.

[88] Noah Wardrip-Fruin. 2007. Three Play Effects–Eliza, Tale-Spin, and Sim City. *Digital Humanities* (2007), 1–2.

[89] William Hollingsworth Whyte. 1980. *The social life of small urban spaces.*

[90] Lauren Winston and Brian Magerko. 2017. Turn-Taking with Improvisational Co-Creative Agents. In *Proceedings of the Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE '17)*. AAAI, Snowbird, UT.

[91] Niels Wouters, John Downs, Mitchell Harrop, Travis Cox, Eduardo Oliveira, Sarah Webber, Frank Vetere, and Andrew Vande Moere. 2016. Uncovering the honeypot effect: How audiences engage with public interactive systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 5–16.

[92] Qian Yang. 2018. Machine Learning as a UX Design Material: How Can We Imagine Beyond Automation, Recommenders, and Reminders?. In *2018 AAAI Spring Symposium Series*.

[93] Georgios N Yannakakis, Antonios Liapis, and Constantine Alexopoulos. 2014. Mixed-initiative co-creativity.. In *FDG*.

[94] Quan-shi Zhang and Song-Chun Zhu. 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 27–39.

[95] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 493–502.